

Ensemble Robustness of Deep Learning Algorithms

Jiashi Feng^{*}, Tom Zahavy[‡], Bingyi Kang^{*}, Huan Xu[†], and Shie Mannor[‡]

^{*}Department of ECE, National University of Singapore

[‡]Department of EE, Technion

[†]Department of ISE, National University of Singapore

Abstract

The question why deep learning algorithms perform so well in practice has attracted increasing research interest. However, most of well-established approaches, such as hypothesis capacity, robustness or sparseness, have not provided complete explanations, due to the high complexity of the deep learning algorithms and their inherent randomness. In this work, we introduce a new approach – ensemble robustness – towards characterizing the generalization performance of generic deep learning algorithms. Ensemble robustness concerns robustness of the *population* of the hypotheses that may be output by a learning algorithm. Through the lens of ensemble robustness, we reveal that a stochastic learning algorithm can generalize well as long as its sensitiveness to adversarial perturbation is bounded in average, or equivalently, the performance variance of the algorithm is small. Quantifying ensemble robustness of various deep learning algorithms may be difficult analytically. However, extensive simulations for seven common deep learning algorithms for different network architectures provide supporting evidence for our claims. Furthermore, our work explains the good performance of several published deep learning algorithms.

1 Introduction

Deep neural networks have been successfully applied in many artificial intelligence tasks and provide state-of-the-art performance. However, the theoretical understandings of learning in deep neural networks has lagged the practical success. The reason why deep learning algorithms perform so well and how to predict their generalization performance are still unclear. Most of well-established approaches, such as hypothesis capacity or sparseness, have not provided complete explanations, due to the high complexity of the deep learning algorithms and their inherent randomness.

Recently, Xu *et al.* [21] established the connection between robustness of an algorithm and its generalization performance. In particular, they point out that if one algorithm is robust (*i.e.*, its empirical loss does not change dramatically for perturbed samples), its generalization performance can also be guaranteed in a positive way. In parallel with that work, several methods based on robustly optimizing the empirical loss are developed for training deep neural network models [17, 22, 15]. However, deep neural networks were shown to be fragile to adversarial perturbation on the inputs [17, 22]. Even an imperceptible perturbation over training samples can corrupt performance of a neural network. Thus the robustness argument in [21] cannot be applied here for explaining the performance of deep learning algorithms. Several researchers also propose to understand deep learning performance through regularization [30, 9], stability [6] and non-convex optimization [2]. Although those works and others are pursuing explanation of the performance of deep neural networks, the reason is still far away from being understood well.

In this work, we present a new approach, *ensemble robustness*, to characterize the generalization performance of deep learning algorithms. Our proposed approach is not intended to give tight performance guarantees for general deep learning algorithms, but rather to pave a way for addressing the question: why deep learning performs so well? Answering this question is difficult, yet we present evidence in

both theory and simulations strongly suggesting that *ensemble robustness* is crucial to the generalization performance of deep learning algorithms.

Ensemble robustness concerns the fact that a randomized algorithm (*e.g.*, many modern deep learning algorithms have inherent randomness) produces a distribution of hypotheses instead of a deterministic one. Therefore, ensemble robustness takes into consideration robustness of the *population* of the hypotheses: even though some hypotheses may be sensitive to perturbation on inputs, an algorithm can still generalize well as long as the most of the hypotheses sampled from the distribution are robust. Ensemble robustness gives a rigorous description to this feature of randomized algorithms.

Through ensemble robustness, we prove that the following holds with a high probability: randomized learning algorithms can generalize well as long as its output hypothesis has bounded sensitiveness to perturbation in average (see Theorem 1). Specified for deep learning algorithms, we reveal that if the hypothesis from different runs of a deep learning method performs consistently well in terms of robustness, the performance of such a deep learning method can be confidently expected. Furthermore, we provide an explanation on the effectiveness of dropout through ensemble robustness.

Although ensemble robustness may be difficult to compute analytically, we demonstrate an empirical estimate of ensemble robustness also works well. We empirically investigate the role of ensemble robustness via extensive simulations on seven common deep learning algorithms. The results provide supporting evidence for our claim: ensemble robustness consistently explain the performance well for deep networks with different architectures and deep learning algorithms. We believe this work could pave a way to understand why deep learning methods perform well in practice.

2 Related Works

Xu *et al.* [21] proposed to consider model robustness for estimating generalization performance for deterministic algorithms, such as for SVM [20] and Lasso [19]. They suggest to use robust optimization to construct learning algorithms, *i.e.*, minimizing the empirical loss with respect to the adversarial perturbed training examples.

Introducing stochasticity to deep learning algorithms has achieved great success in practice and also receives theoretical investigation. Hardt *et al.* [6] analyzed the stability property of Stochastic Gradient Descent (SGD) methods, while dropout training was introduced as a way to control over-fitting by randomly omitting subsets of features at each iteration of a training procedure [10, 30]. Recently, Baldi and Wager *et al.* [1, 18] tried to explain why dropout helps in avoiding over-fitting as a regularization method, while Gal *et al.* [3] explains it as a Bayesian approximation for a Gaussian process. Different from those works, this work will extend the results in [21] to randomized algorithms in order to analyze them from an ensemble robustness perspective.

Adversarial examples for deep neural networks were first introduced in [17], while some recent works propose to utilize them as a regularization technique for training deep models [22, 5, 15]. However, all of those works attempt to find the “worst case” examples in a local neighborhood of the original training data and therefore may be seen as an approximation to the robust optimization approach developed in [21].

Analyzing generalization and robustness properties of deep networks has also gained attention from the deep reinforcement learning community. For instance, Mnih *et al.* [13] trained deep networks to play Atari2600 games, and demonstrated that the same algorithm and network architecture can generalize well for different environments. In order to choose the network architecture they evaluated the policy with an epsilon greedy scheme, thus assuring that the policy is robust to small state perturbations. Levine *et al.* [12] used deep networks to learn control policies for robots. They suggested to apply trajectory optimization methods on sampled robot trajectories and used a deep network to generalize from the samples. To evaluate robustness of the network they perturbed the target object at each training trial, and used varying amounts of perturbation while testing. Heess *et al.* [7] suggested to solve continuous control policies with value gradient algorithm that uses deep networks as a function approximator. They showed that stochastic algorithms are more robust than their deterministic counterparts by evaluating

the algorithm on a variety of network architectures.

3 Preliminaries

In this work, we investigate the generalization property of learning algorithms in deep neural networks, by establishing their PAC bounds. In this section, we provide some preliminary facts that are necessary for developing the approach of ensemble robustness. After introducing the problem setting we are interested in, we in particular highlight the inherent randomness of deep learning algorithms and give a formal description on randomized learning algorithms. Then, we briefly review the relation between robustness and generalization performance established in [21].

Problem setup We here introduce the learning setup for deep neural networks, which follows a standard one for supervised learning. More concretely, we have \mathcal{Z} and \mathcal{H} as the sample set and the hypothesis set respectively. The training sample set $\mathbf{s} = \{s_1, \dots, s_n\}$ consists of n i.i.d. samples generated by an unknown distribution μ , and the target of learning is to obtain a neural network that minimizes *expected* classification error over the i.i.d. samples from μ . Throughout the paper, we consider the training set \mathbf{s} with a fixed size of n .

We denote the learning algorithm as \mathcal{A} , which is a mapping from \mathcal{Z}^n to \mathcal{H} . We use $h_{\mathbf{s}} \leftarrow \mathcal{A} : \mathbf{s}$ to denote the learned hypothesis given the training set \mathbf{s} . We consider the loss function $\ell(h, z)$ whose value is nonnegative and upper bounded by M . Let $\mathcal{L}(\cdot)$ and $\ell_{\text{emp}}(\cdot)$ denote the expected error and the training error for a learned hypothesis $h_{\mathbf{s}}$, i.e.,

$$\mathcal{L}(h_{\mathbf{s}}) \triangleq \mathbb{E}_{z \sim \mu} \ell(h_{\mathbf{s}}, z), \text{ and } \ell_{\text{emp}}(h_{\mathbf{s}}) \triangleq \frac{1}{n} \sum_{s_i \in \mathbf{s}} \ell(h_{\mathbf{s}}, s_i). \quad (1)$$

We are going to characterize the generalization error $|\mathcal{L}(h_{\mathbf{s}}) - \ell_{\text{emp}}(h_{\mathbf{s}})|$ of deep learning algorithms in the following section.

Randomized algorithms Most of modern deep learning algorithms are in essence randomized ones, which in fact map a training set \mathbf{s} to a *distribution* of hypotheses $\Delta(\mathcal{H})$ instead of a single hypothesis. For example, running a deep learning algorithm \mathcal{A} with dropout for multiple times will produce different hypotheses which can be deemed as samples from the distribution $\Delta(\mathcal{H})$. This is an important observation we make for deep learning analysis in this work, and we will point out such randomness actually plays an important role for deep learning algorithms to perform well. Therefore, before proceeding to analyze the performance of deep learning, we provide a formal definition of randomized learning algorithms here.

Definition 1 (Randomized Algorithms). *A randomized learning algorithm \mathcal{A} is a function from \mathcal{Z}^n to a set of distributions of hypotheses $\Delta(\mathcal{H})$, which outputs a hypothesis $h_{\mathbf{s}} \sim \Delta(\mathcal{H})$ with a probability $\pi_{\mathbf{s}}(h)$.*

When learning with a randomized algorithm, the target is to minimize the expected empirical loss for a specific output hypothesis $h_{\mathbf{s}}$, similar to the ones in (1). Here ℓ is the loss incurred by a specific output hypothesis by one instantiation of the randomized algorithm \mathcal{A} .

Examples of the internal randomness of a deep learning algorithm \mathcal{A} include dropout rate (the parameter for a Bernoulli distribution for randomly masking certain neurons), random shuffle among training samples in stochastic gradient descent, the initialization of weights for different layers, to name a few.

Robustness and generalization Xu *et al.* [21] for the first time established the relation between algorithmic robustness and generalization. An algorithm is robust if following holds: if two samples are close to each other, their associated losses are also close. For being self-contained, we here briefly review the algorithmic robustness and its induced generalization guarantee.

Definition 2 (Robustness [21]). *Algorithm \mathcal{A} is $(K, \epsilon(\cdot))$ robust, for $K \in \mathbb{N}$ and $\epsilon(\cdot) : \mathcal{Z}^n \rightarrow \mathbb{R}$, if \mathcal{Z} can be partitioned into K disjoint sets, denoted by $\{C_i\}_{i=1}^K$, such that the following holds for all $\mathbf{s} \in \mathcal{Z}^n$:*

$$\begin{aligned} &\forall \mathbf{s} \in \mathbf{s}, \forall z \in \mathcal{Z}, \forall i = 1, \dots, K : \\ &\text{if } s, z \in C_i, \text{ then } |\ell(\mathcal{A}_{\mathbf{s}}, s) - \ell(\mathcal{A}_{\mathbf{s}}, z)| \leq \epsilon(n). \end{aligned}$$

Based on the above robustness property of algorithms, Xu *et al.* [21] prove that a robust algorithm also generalizes well. Motivated by their results, Shaham *et al.* [15] proposed adversarial training algorithm to minimize the empirical loss over synthesized adversarial examples. However, those results cannot be applied for characterizing the performance of modern deep learning models well.

4 Ensemble Robustness

In order to explain the good performance of deep learning, one needs to understand the internal randomness of deep learning algorithms and the population performance of the multiple possible hypotheses. Intuitively, a single output hypothesis cannot be robust to adversarial perturbation on training samples and the deterministic robustness argument in [21] cannot be applied here. Fortunately, deep learning algorithms generally output the hypothesis sampled from a distribution of hypotheses. Therefore, even some samples are not “nice” for one specific hypothesis, with a large probability it cannot fail *most* of the hypotheses from the produced distribution. Thus, deep learning algorithms are able to generalize well. Such intuition motivates us to introduce the concept of *ensemble robustness* that is defined over a distribution of hypotheses output by a deep learning algorithm.

Definition 3 (Ensemble Robustness). *A randomized algorithm \mathcal{A} is $(K, \bar{\epsilon}(n))$ ensemble robust, for $K \in \mathbb{N}$ and $\bar{\epsilon}(n)$, if \mathcal{Z} can be partitioned into K disjoint sets, denoted by $\{C_i\}_{i=1}^K$, such that the following holds for all $\mathbf{s} \in \mathcal{Z}^n$:*

$$\begin{aligned} &\forall \mathbf{s} \in \mathbf{s}, \forall i = 1, \dots, K : \\ &\text{if } s \in C_i, \text{ then } \mathbb{E}_{\mathcal{A}} \max_{z \in C_i} |\ell(\mathcal{A}_{\mathbf{s}}, s) - \ell(\mathcal{A}_{\mathbf{s}}, z)| \leq \bar{\epsilon}(n). \end{aligned}$$

Here the expectation is taken w.r.t. the internal randomness of the algorithm \mathcal{A} .

Ensemble robustness is a “weaker” requirement for the model compared with the robustness proposed in [21] and it fits better for explaining deep learning. In the following section, we demonstrate through simulations that a deep learning model is not robust but it is indeed ensemble robust. So in practice the deep model can still achieve good generalization performance.

An algorithm with strong ensemble robustness can provide good generalization performance in expectation w.r.t. the generated hypothesis, as stated in the following theorem.

Theorem 1. *Let \mathcal{A} be a randomized algorithm with $(K, \bar{\epsilon}(n))$ ensemble robustness over the training set \mathbf{s} , with $|\mathbf{s}| = n$. Let $\Delta(\mathcal{H}) \leftarrow \mathcal{A} : \mathbf{s}$ denote the output hypothesis distribution of \mathcal{A} . Then for any $\delta > 0$, with probability at least $1 - \delta$ with respect to the random draw of the \mathbf{s} and $h \sim \Delta(\mathcal{H})$, the following holds:*

$$|\mathcal{L}(h) - \ell_{\text{emp}}(h)| \leq \sqrt{\frac{nM\bar{\epsilon}(n) + 2M^2}{\delta n}}.$$

Note that in the above theorem, we hide the dependency of the generalization bound on K in ensemble robustness measure $\bar{\epsilon}(n)$. Due to space limitations, all the technical lemmas and details of the proofs throughout the paper are deferred to supplementary material. Theorem 1 leads to following corollary which gives a way to minimize expected loss directly.

Corollary 1. Let \mathcal{A} be a randomized algorithm with $(K, \bar{\epsilon}(n))$ ensemble robustness. Let C_1, \dots, C_K be a partition of \mathcal{Z} , and write $z_1 \sim z_2$ if z_1, z_2 fall into the same C_k . If the training sample \mathbf{s} is generated by i.i.d. draws from μ , then with probability at least $1 - \delta$, the following holds over $h \in \mathcal{H}$

$$\mathcal{L}(h) \leq \frac{1}{n} \sum_{i=1}^n \max_{z_i \sim s_i} \ell(h, z_i) + \sqrt{\frac{nM\bar{\epsilon}(n) + 2M^2}{\delta n}}.$$

Corollary 1 suggests that one can minimize the expected error of a deep learning algorithm effectively through minimizing the empirical error over the training samples s_i perturbed in an adversarial way. In fact, such an adversarial training strategy has been exploited in [22, 15].

Theorem 2. Let \mathcal{A} be a randomized algorithm with $(K, \bar{\epsilon}(n))$ ensemble robustness over the training set \mathbf{s} , where $|\mathbf{s}| = n$. Let $\Delta(\mathcal{H})$ denote the output hypothesis distribution of the algorithm \mathcal{A} on the training set \mathbf{s} . Suppose following variance bound holds:

$$\text{var}_{\mathcal{A}} \left[\max_{z \sim s_i} |\ell(\mathcal{A}_{\mathbf{s}}, s_i) - \ell(\mathcal{A}_{\mathbf{s}}, z)| \right] \leq \alpha$$

Then for any $\delta > 0$, with probability at least $1 - \delta$ with respect to the random draw of the \mathbf{s} and $h \sim \Delta(\mathcal{H})$, we have

$$|\mathcal{L}(\mathcal{A}_{\mathbf{s}}) - \ell_{\text{emp}}(\mathcal{A}_{\mathbf{s}})| \leq \bar{\epsilon}(n) + \frac{1}{\sqrt{2\delta}}\alpha + M\sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}}$$

Theorem 2 suggests that controlling the variance of the deep learning model can substantially improve the generalization performance. Note that here we need to consider the trade-off between the expectation and variance of ensemble robustness. To see this, consider following two extreme examples. When $\alpha = 0$, we do not allow any variance in the output of the algorithm \mathcal{A} . Thus, \mathcal{A} reduces to a deterministic one. To achieve the above upper bound, it is required that the output hypothesis satisfies $\max_{z \in C_i} |\ell(h, s_i) - \ell(h, z)| \leq \epsilon(n)$. However, due to the intriguing property of deep neural networks [17], the deterministic model robustness measure $\epsilon(n)$ (ref. Definition 2) is usually large. In contrast, when the hypotheses variance α can be large enough, there are multiple possible output hypotheses from the distribution $\Delta(\mathcal{H})$. We fix the partition of \mathcal{Z} as C_1, \dots, C_K . Then,

$$\begin{aligned} & \mathbb{E}_{\mathcal{A}} \left[\max_{z \sim \mathbf{s} \in C_i} |\ell(h, s) - \ell(h, z)| \right] \\ &= \sum_{j \in \Delta(\mathcal{H})} \mathbb{P}\{h = h_j\} \max_{z \sim \mathbf{s} \in C_i} |\ell(h_j, s) - \ell(h_j, z)| \\ &\leq \sum_{j \in \Delta(\mathcal{H})} \mathbb{P}\{h = h_j\} \max_{z \sim \mathbf{s} \in C_i} \max_{h \in \Delta(\mathcal{H})} |\ell(h, s) - \ell(h, z)| \\ &\leq \max_{z \sim \mathbf{s} \in C_i} \max_{h \in \Delta(\mathcal{H})} |\ell(h, s) - \ell(h, z)|. \end{aligned}$$

Therefore, allowing certain variance on produced hypotheses, a randomized algorithm can tolerate the non-robustness of some hypotheses to certain samples. As long as the ensemble robustness is small, the algorithm can still perform well.

5 Understanding Dropout via Ensemble Robustness

In this section, we illustrate how ensemble robustness can well characterize the performance of various training strategies of deep learning. In particular, we take the dropout as a concrete example.

Dropout is a widely used technique for optimizing deep neural network models. We demonstrate that dropout is a random scheme to perturb the algorithm. During dropout, at each step, a random fraction of the units are masked out in a round of parameter updating.

Assumption 1. We assume the randomness of the algorithm \mathcal{A} is parametrized by $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_L) \in \mathcal{R}$ where $\mathbf{r}_l, l = 1, \dots, L$ are random elements drawn independently.

For a deep neural network consisting of L layers, the random variable \mathbf{r}_l is the dropout randomness for the l -th layer. The next theorem establishes the generalization performance for the neural network with dropout training.

Theorem 3 (Generalization of Dropout Training). *Consider an L -layer neural network trained by dropout. Let \mathcal{A} be an algorithm with $(K, \bar{\epsilon}(n))$ ensemble robustness. Let $\Delta(\mathcal{H})$ denote the output hypothesis distribution of the randomized algorithm \mathcal{A} on a training set \mathbf{s} . Assume there exists a $\beta > 0$ such that,*

$$\sup_{\mathbf{r}, \mathbf{t}} \sup_{z \in \mathcal{Z}} |\ell(\mathcal{A}_{\mathbf{s}, \mathbf{r}}, z) - \ell(\mathcal{A}_{\mathbf{s}, \mathbf{t}}, z)| \leq \beta \leq L^{-3/4},$$

with \mathbf{r} and \mathbf{t} only differing in one element. Then for any $\delta > 0$, with probability at least $1 - \delta$ with respect to the random draw of the \mathbf{s} and $h \sim \Delta(\mathcal{H})$,

$$\mathcal{L}(h_{\mathbf{s}, \mathbf{r}}) - \ell_{\text{emp}}(h_{\mathbf{s}, \mathbf{r}}) \leq \bar{\epsilon}(n) + \sqrt{2 \log(1/\delta)/L} + \sqrt{\frac{2K \ln 2 + 2 \ln(2/\delta)}{n}}.$$

Theorem 3 also establishes the relation between the depth of a neural network model and the generalization performance. It suggests that when using dropout training, controlling the variance β of the empirical performance over different runs is important: when β converges at the rate of $L^{-3/4}$, increasing the layer number L will improve the performance of a deep neural network model. However, simply making L larger without controlling β does not help. Therefore, in practice, we usually use voting from multiple models to reduce the variance and thus decrease the generalization error [8]. Also, when dropout training is applied for more layers in a neural network model, smaller variance of the model performance is preferred. This can be compensated by increasing the size of training examples or ensemble of multiple models.

6 Simulations

This section is devoted to simulations for quantitatively and qualitatively demonstrating how ensemble robustness of a deep learning method explains its performance. We first introduce our experiment settings and implementation details.

6.1 Experiment Settings

Dataset We conduct simulations on the benchmark MNIST handwritten digit dataset which contains 60,000 training samples and 10,000 test samples [11]. The digit images have been size-normalized and centered in pre-processing.

Network architecture Without explicit explanation, we use a convolutional network with the following architecture throughout the simulations. The network contains two convolutional layers (with 32 and 64 5×5 filters respectively), each of which is followed by a max-pooling layer (3×3 and 2×2 respectively), and two fully connected layers (with 200 and 10 units) on top. The output of the last fully-connected layer is fed to a 10-way softmax and cross-entropy loss function. The magnitude of the loss is upper bounded by 1. We call this network as *Network-I*. In order to avoid the bias brought by specific network architecture on our observations, we also consider another two alternative architectures, called *Network-II* and *Network-III* respectively. Compared with *Network-I*, *Network-II* only has one convolutional layer (with 64 filters) and *Network-III* has an additional fully connected layer (with 100 units) on top of the 200-unit layer. In short, the three networks are ranked as follows according to their complexity: *Network-III* > *Network-I* > *Network-II*.

Table 1: Comparison between deterministic and randomized deep learning methods on their training error, test error and misclassification rate on testing examples with perturbation.

Network	deterministic			randomized		
	training	test	perturbed	training	test	perturbed
N100-100-10	0.57%	0.84%	2.1%	0.56%	0.85%	0.86%
N200-200-10	0.20%	0.71%	1.94%	0.20%	0.70%	0.87%
N100-100-100-10	0%	0.51%	1.68%	0%	0.49%	0.45%
N200-200-200-10	0%	0.38%	1.54%	0%	0.38%	0.35%

Parameter setting We trained this network using a mini-batch of 100 training examples at a time, a learning rate of 0.01, a learning rate decay of 1×10^{-6} and a momentum of 0.9. For the Dropout algorithm we add a dropout layer with a dropout rate of 0.5 after the first fully connected layer.

Compared algorithms We evaluate and compare ensemble robustness as well as the generalization performance for following 7 deep learning algorithms: stochastic gradient descent (SGD), SGD plus dropout, Prioritized Sweeping (PS) applied for supervised learning [14], PS plus dropout, and three versions of adversarial training algorithms used in [17, 22, 15], where the magnitude of perturbation is measured by its ℓ_1 , ℓ_2 and ℓ_∞ norm respectively.

6.2 Robustness vs. Ensemble Robustness

We first present simulations that motivate the concept of ensemble robustness: why a deep neural network that is not robust to adversarial sample perturbation can perform well in practice? In the simulations, we examine four network models with only fully connected layers: two of the examined networks have two fully connected layers and the other two have three fully connected layers. The architecture configuration is inspired by [17] and the layer sizes for the used networks are provided in Table 1. We generate a set of adversarial instances for a given network, using the approach proposed by [17], and feed these examples for the corresponding network to evaluate its classification error. We in particular compare the performance between a *deterministic* hypothesis and a *randomly* sampled hypothesis. The randomized hypothesis is produced by following the definition of ensemble robustness: we run the network training for 10 times with different initialization, and randomly sample one hypothesis from the pool of 10 hypothesis for testing. Note that we do not perform any multiple model combination here. The results are given in Table 1. From the table, one can make the following observation. The deterministic hypothesis performs as well as its randomized counterpart on training and normal testing examples (without adversarial perturbation). However, on perturbed samples, the deterministic one performs much worse than the randomized one. This suggests that deterministic robustness argument (as the one proposed by [21]) is not applicable for deep learning algorithms: the algorithm is not robust to adversarial perturbation but indeed performs well in practice. Instead, we should consider ensemble robustness w.r.t. a population of hypothesis, as the randomized hypothesis performs well on specifically designed adversarial samples.

6.3 Empirical Ensemble Robustness vs. Generalization

We then proceed to investigate the relation between ensemble robustness and generalization performance of a deep learning algorithm to validate the results in Theorem 1. As ensemble robustness involves taking expectation over all the possible output hypothesis, it is computationally difficult to exactly measure ensemble robustness for different deep learning algorithms. In this simulation, we take the empirical average of robustness to adversarial perturbation from 30 different hypotheses of the same learning algorithm as its ensemble robustness.

In particular, we aim to empirically demonstrate that a deep learning algorithm with stronger ensemble robustness presents better generalization performance (Theorem 1). Recall the definition of ensemble

robustness in Definition 3, another obstacle in calculating ensemble robustness is to find the most adversarial perturbation Δs (or equivalently the most adversarial example $z = s + \Delta s$) for a specific training sample $s \in \mathbf{s}$ within a partition set C_i . We therefore employ an approximate search strategy for finding the adversarial examples. More concretely, we optimize the following first-order Taylor expansion of the loss function as a surrogate for finding the adversarial example :

$$\Delta s_i \in \arg \max_{\|\Delta s_i\| \leq r} \ell(s_i) + \langle \nabla \ell_{s_i}(s), \Delta s_i \rangle, \quad (2)$$

with a pre-defined magnitude constraint r on the perturbation Δs_i . In the simulations, we vary the magnitude r in order to calculate the empirical ensemble robustness at different perturbation levels. More details about this approximation technique can be found in [15].




Figure 1: Empirical ensemble robustness $\bar{\epsilon}_{\text{emp}}$ for different deep learning algorithms with Network I. Here the adversarial perturbation is given with different magnitudes (x-axis).

We optimize the networks with different architectures depicted in Section 6.1 using seven algorithms: vanilla SGD, vanilla SGD plus dropout, prioritized sweeping (PS), PS plus dropout, and three versions of the adversarial training methods used in [17, 22, 15] with different metric on the magnitude of perturbation r . Each deep learning algorithm is run for 10 times with different random seeds (including the ones for initialization, training example shuffle and dropout).

We then calculate the *empirical* ensemble robustness by averaging the difference between the loss of the algorithm on the training samples and the adversarial samples output by the method in (2):

$$\bar{\epsilon}_{\text{emp}} = \frac{1}{T} \sum_{t=1}^T \max_{i \in \{1, \dots, n\}} |\ell(\mathcal{A}_s^{(t)}, s_i) - \ell(\mathcal{A}_s^{(t)}, s_i + \Delta s_i)|,$$

with $T = 10$ denoting the number of different runs of the algorithm in the simulations. We monotonically increase the magnitude r of the perturbation Δs_i in (2) to generate adversarial examples at different noise levels.

The performance of different algorithms and different networks on test set are given in Table 2. We plot the empirical ensemble robustness of different deep learning algorithms $\bar{\epsilon}$ against magnitude of the

sample perturbation in Figure 1. We do not plot the curves for PS and PS+dropout algorithms as they overlap with the ones of SGD and SGD+dropout heavily, respectively.

One can observe that our simulations have yielded positive evidence on MNIST to support our claim on the relation between ensemble robustness and algorithm generalization performance in Theorem 1. Comparing the results given in Table 2 and Figure 1, one can observe that adversarial training based methods consistently present stronger ensemble robustness (smaller $\bar{\epsilon}$) than pure SGD and dropout, across different magnitude of the perturbation. Also, they produce lower testing error than SGD and dropout. This verifies our proposed ensemble robustness indeed indicates the performance of a deep learning method accurately. We also compute ensemble robustness for different networks (*i.e.*, Network I, II and III) in the supplementary material due to space limitation.

Table 2: Classification error (in %) of different learning algorithms with different networks on the MNIST dataset.

Alg.	Network-I	Network-II	Network-III
Vanilla SGD	0.89	3.77	2.19
Prioritized Sweeping	0.82	3.52	2.93
SGD + Dropout	0.78	2.39	1.08
PS + Dropout	0.74	2.06	1.89
ℓ_1 adversarial	0.67	1.09	0.7
ℓ_2 adversarial	0.66	1.96	0.85
ℓ_∞ adversarial	0.65	1.88	0.7

7 Conclusions

We introduced ensemble robustness to analyze the performance of randomized algorithms, in particular deep learning algorithms. We established both theoretically and experimentally that ensemble robustness of an algorithm indicates its generalization performance well. We also demonstrated why dropout works well through analyzing its ensemble robustness. Experiments on real world dataset also confirmed our claims. In this work, we do not target at establishing very tight performance bounds. Instead, we provide results that give insight on existing deep learning architectures which may lead to new algorithms and architectures in the future.

References

- [1] Pierre Baldi and Peter J Sadowski. Understanding dropout. In *Advances in Neural Information Processing Systems*, 2013.
- [2] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The loss surface of multilayer networks. In *International Conference on Artificial Intelligence and Statistics*, 2015.
- [3] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Insights and applications. In *Deep Learning Workshop, ICML*, 2015.
- [4] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- [5] Shixiang Gu and Luca Rigazio. Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068*, 2014.
- [6] Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.

- [7] Nicolas Heess, Gregory Wayne, David Silver, Tim Lillicrap, Tom Erez, and Yuval Tassa. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, pages 2926–2934, 2015.
- [8] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2014.
- [9] Prateek Jain, Vivek Kulkarni, Abhradeep Thakurta, and Oliver Williams. To drop or not to drop: Robustness, consistency and differential privacy properties of dropout. *arXiv preprint arXiv:1503.02031*, 2015.
- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [11] Yann LeCun, Corinna Cortes, and Christopher JC Burges. The mnist database of handwritten digits, 1998.
- [12] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 2016.
- [13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [14] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *International Conference on Learning Representations*, 2016.
- [15] Uri Shaham, Yutaro Yamada, and Sahand Negahban. Understanding adversarial training: Increasing local stability of neural nets through robust optimization. *arXiv preprint arXiv:1511.05432*, 2015.
- [16] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [17] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- [18] Stefan Wager, Sida Wang, and Percy S Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems*, pages 351–359, 2013.
- [19] Huan Xu, Constantine Caramanis, and Shie Mannor. Robust regression and lasso. In *Advances in Neural Information Processing Systems*, 2009.
- [20] Huan Xu, Constantine Caramanis, and Shie Mannor. Robustness and regularization of support vector machines. *The Journal of Machine Learning Research*, 10:1485–1510, 2009.
- [21] Huan Xu and Shie Mannor. Robustness and generalization. *Machine learning*, 86(3):391–423, 2012.
- [22] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- [23] Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589, 2014.

- [24] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, 2013.
- [25] Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- [26] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing by virtual adversarial examples. In *International Conference on Learning Representation*, 2016.
- [27] Nikolaos Pitelis, Chris Russell, and Lourdes Agapito. Semi-supervised learning using an unsupervised atlas. In *Machine Learning and Knowledge Discovery in Databases*, pages 565–580. Springer, 2014.
- [28] Antti Rasmus, Mathias Berglund, Mikko Honkela, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pages 3532–3540, 2015.
- [29] Salah Rifai, Yann N Dauphin, Pascal Vincent, Yoshua Bengio, and Xavier Muller. The manifold tangent classifier. In *Advances in Neural Information Processing Systems*, pages 2294–2302, 2011.
- [30] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [31] Steven C Sudderth and YL Kergosien. Rule-injection hints as a means of improving network performance and learning time. In *Neural Networks*, pages 120–129. Springer, 1990.
- [32] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.
- [33] Jason Weston, Frédéric Ratle, Hossein Mobahi, and Ronan Collobert. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*, pages 639–655. Springer, 2012.

A Additional Simulations

In this section, we provide additional simulations to empirically validate our theories about ensemble robustness. First, we conduct simulations to investigate the relation between performance variation and generalization as stated in Theorem 2.

A.1 Performance Variance vs. Generalization

We now investigate whether the performance (or more concretely, the variance of empirical ensemble robustness) variance from multiple runs of the deep learning algorithms also affects the generalization performance, as stated in Theorem 2. To empirically estimate the variance of $\max_{z \sim s} |\ell(\mathcal{A}_s, s) - \ell(\mathcal{A}_s, z)|$, we run all the deep learning algorithms for 30 times with Network II. The simulations on Network I and III also present similar trend. Here the adversarial perturbed examples $z = s + \Delta s$ are also provided by the strategy in (2). The variance statistics as well as the final performance of different algorithms on MNIST dataset are presented in Figure 2. One can observe from the table that controlling the performance variance is important for deep learning algorithms. Also, this gives a computationally feasible way to estimate the generalization performance of deep learning methods.




Figure 2: Variance of hypothesis robustness and final performance of different deep learning algorithms with Network II. \square indicates the classification accuracy and the error bar describes the variance. It can be observed that algorithm with smaller robustness variance also gives more accurate classification results.

A.2 Visualization on Learned Representations

This section is devoted to providing intuitive understanding on the implications of ensemble robustness for different deep learning algorithms. We employ tSNE [32] to visualize the learned representations from Network-I on MNIST in Figure 3 from five different learning algorithms. We also visualize the raw data with color codes for different categories. To demonstrate the distribution of different types of samples, we encode the “training”, “test” and “adversarial” ones by different shapes. We use different colors to indicate whether they are classified correctly. Here the adversarial examples are generated with relatively large perturbation magnitude 0.4.

The visualization results offer following observations. First of all, compared with raw data, representations output from all the deep learning algorithms show clearer cluster (sub-manifold) structure. Deep learning algorithms are effective at extracting discriminative representations. Secondly, from the visualization of SGD and dropout results, one can observe that most of the perturbed examples are misclassified (there are many red crosses). This indicates these two algorithms possess ensemble robustness, and their generalization performance is also worse than others. In contrast, for the learning algorithms with explicit adversarial training, much fewer perturbed examples are misclassified. These algorithms have stronger ensemble robustness and better generalization performance. This again confirms our claim on the importance of ensemble robustness for algorithms generalization performance. Thirdly, taking a closer look at (d)-(f), we can observe that representations of adversarial examples are very close to the original examples. Recalling the dependency of ensemble robustness on the sample space partitions C_i , this implies the algorithms in (d)-(f) are associated with a larger C_i capable of accommodating more perturbed samples (and correspondingly smaller K). This feature also endows the algorithms with better generalization performance.

Table 3: A collection of previously reported MNIST test errors (in %) followed by the results with our proposed semi-supervised learning method. Standard deviation in parentheses.

# of used labels	100	1000	All
Embedding [33]	16.86	5.73	1.50
Transductive SVM [33]	16.81	5.38	1.40
MTC [29]	12.03	3.64	0.81
Pseudo-label [24]	10.49	3.46	
AtlasRBF [27]	8.10(± 0.95)	3.68(± 0.12)	1.31
DGN [23]	3.33(± 0.14)	2.40(± 0.02)	0.96
DBM, Dropout [30]			0.79
Adversarial [22]			0.78
Virtual Adversarial [26]	2.12	1.32	0.64(± 0.03)
Ladder Network [28]	1.06(± 0.37)	0.84(± 0.08)	0.57(± 0.02)
Ours	1.03 (± 0.32)	0.81 (± 0.10)	0.55 (± 0.02)

A.3 Application for Semi-supervised Learning

In this subsection, we showcase a new semi-supervised learning method for deep neural network training, inspired by the result given in Corollary 1. The idea of using unsupervised learning to complement supervision is not new. Combining an auxiliary task to help train a neural network was investigated in [31]. By sharing the hidden representations among more than one task, the network generalizes better. There are multiple choices for the unsupervised task. Here, we follow Corollary 1 and define the unsupervised learning task as the one to classify the slightly perturbed training samples (without labels) into the same class as original samples without perturbation. With the experiments on the MNIST dataset, we compare our method to other semi-supervised methods to demonstrate how the introduced perturbation on training samples could help improve the classification accuracy, without requiring additional training samples. The network architecture and parameter setting in this experiment also follow the ones introduced in Section 6.1.

For evaluating semi-supervised learning, we use the standard 10,000 test samples as a held-out test set and randomly split the standard 60,000 training samples into a 10,000-sample validation set and use 50,000 as the training set. From the training set, we randomly choose 100, 1000, or all labels for the supervised cost. We also balance the classes to ensure that no particular class is over-represented. We repeat each training 10 times, varying the random seed that is used for the splits. The hyperparameter we tune for our proposed method is the magnitude of adversarial perturbation on the unlabeled training examples. We optimize such a hyperparameter with a search grid $\{0.01, 0.1, 0.2, 0.5\}$. The results presented in Table 3 show that the proposed method outperforms all the previously reported results.

A.4 Ensemble Robustness of Different Networks

We plot the empirical ensemble robustness of different learning algorithms (five in total) and different networks (*i.e.*, Network I, II and III) in Figure 4. Comparing this results with the numbers shown in Table 2 confirms the consistency between the ensemble robustness and generalization performance.

B Technical Lemmas

Lemma 1. *For a randomized learning algorithm \mathcal{A} with $(K, \bar{\epsilon}(n))$ uniform ensemble robustness, and loss function ℓ such that $0 \leq \ell(h, z) \leq M$, we have,*

$$\mathbb{P}_{\mathbf{s}} \left\{ \mathbb{E}_{\mathcal{A}} |\mathcal{L}(h) - \ell_{\text{emp}}(h)| \leq \bar{\epsilon}(n) + M \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}} \right\} \geq 1 - \delta,$$

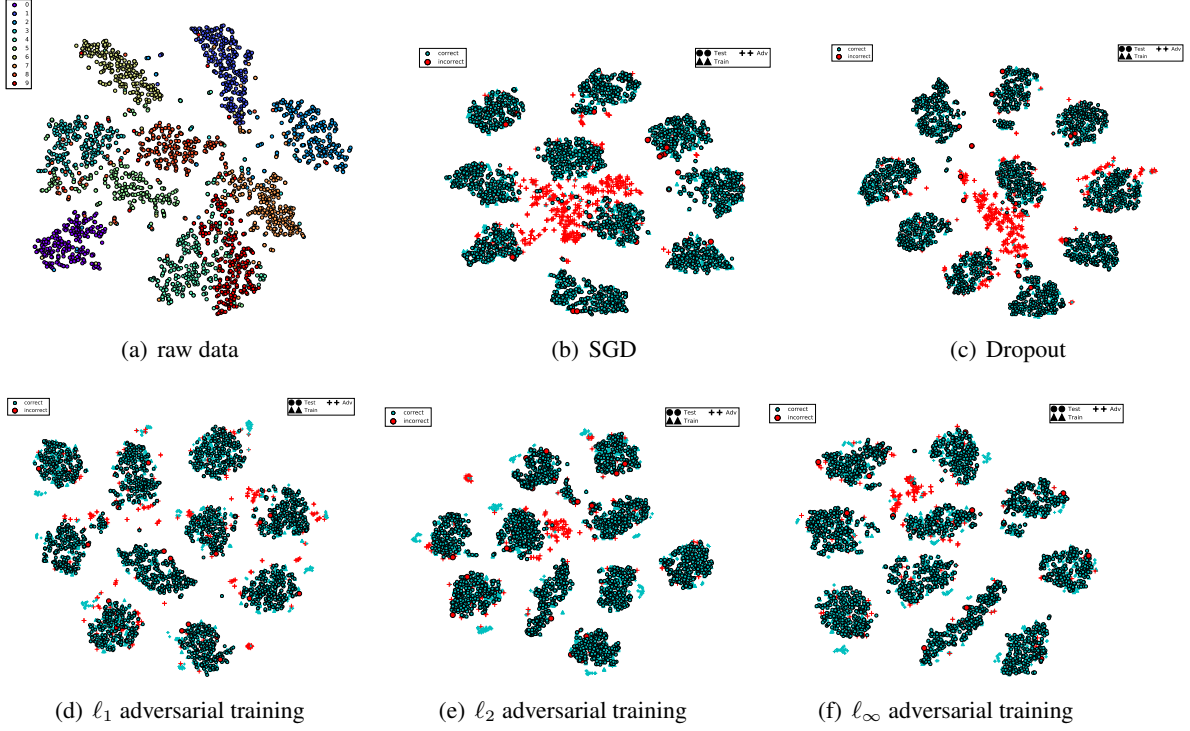


Figure 3: tSNE visualization on the learned representation from a deep convolutional neural network trained with different algorithms (on the same training dataset). (a) raw data; (b) vanilla SGD; (c) SGD plus dropout; (d)-(f) adversarial training with different measures on the magnitude of perturbation (ref. (2)). Different types of the sample are marked by different shapes: \triangle for training, \circ for test and $+$ for perturbed examples. In (a) samples are colored by their groundtruth categories. In (b)-(f), samples classified correctly are marked by green and those not correctly classified are marked by red color. Similar to (a), each cluster of samples corresponds to one category. Compared with SGD and dropout, the adversarial training algorithms in (d)-(f) provides more compact clusters with larger margin among different clusters. In addition, their classification accuracy on perturbed samples is much higher than SGD and dropout.

where we use $\mathbb{P}_{\mathbf{s}}$ to denote the probability w.r.t. the choice of \mathbf{s} , and $|\mathbf{s}| = n$.

Proof. Given a random choice of training set \mathbf{s} with cardinality of n , let N_i be the set of index of points of \mathbf{s} that fall into the C_i . Note that $(|N_1|, \dots, |N_K|)$ is an i.i.d. multinomial random variable with parameters n and $(\mu(C_1), \dots, \mu(C_K))$. The following holds by the Bretaganolle-Huber-Carol inequality:

$$\mathbb{P}_{\mathbf{s}} \left\{ \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| \geq \lambda \right\} \leq 2^K \exp \left(\frac{-n\lambda^2}{2} \right).$$

We have




Figure 4: Empirical ensemble robustness of different learning algorithms as well as different networks. The perturbation magnitude is controlled within 0.4.

$$\begin{aligned}
& \mathbb{E}_{\mathcal{A}} |\mathcal{L}(\mathcal{A}_{\mathbf{s}}) - \ell_{\text{emp}}(\mathcal{A}_{\mathbf{s}})| \\
&= \mathbb{E}_{\mathcal{A}} \left| \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \mu(C_i) - \frac{1}{n} \sum_{i=1}^n \ell(\mathcal{A}_{\mathbf{s}}, s_i) \right| \\
&\stackrel{(a)}{\leq} \mathbb{E}_{\mathcal{A}} \left| \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \frac{|N_i|}{n} - \frac{1}{n} \sum_{i=1}^n \ell(\mathcal{A}_{\mathbf{s}}, s_i) \right| \\
&\quad + \mathbb{E}_{\mathcal{A}} \left| \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \mu(C_i) - \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \frac{|N_i|}{n} \right| \\
&\stackrel{(b)}{\leq} \sum_{i=1}^K \mathbb{E}_{\mathcal{A}} \left| \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \frac{|N_i|}{n} - \frac{1}{n} \sum_{j \in N_i} \ell(\mathcal{A}_{\mathbf{s}}, s_j) \right| \\
&\quad + \mathbb{E}_{\mathcal{A}} \left| \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \mu(C_i) - \sum_{i=1}^K \mathbb{E}_{z \sim \mu}(\ell(\mathcal{A}_{\mathbf{s}}, z) | z \in C_i) \frac{|N_i|}{n} \right| \\
&\leq \frac{1}{n} \sum_{i=1}^K \sum_{j \in N_i} \mathbb{E}_{\mathcal{A}} \left(\max_{z \in C_i} |\ell(\mathcal{A}_{\mathbf{s}}, s_j) - \ell(\mathcal{A}_{\mathbf{s}}, z)| \right) + \max_{z \in \mathcal{Z}} |\ell(\mathcal{A}_{\mathbf{s}}, z)| \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| \quad (3) \\
&\stackrel{(c)}{\leq} \bar{\epsilon}(n) + M \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| \\
&\stackrel{(d)}{\leq} \bar{\epsilon}(n) + M \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}} \quad (4)
\end{aligned}$$

Here the inequalities (a) and (b) are due to triangle inequality, (c) is from the definition of ensemble robustness and the fact that the loss function is upper bounded by M , and (d) holds with a probability greater than $1 - \delta$. \square

Lemma 2. For a randomized learning algorithm \mathcal{A} with $(K, \bar{\epsilon}(n))$ uniform ensemble robustness, and loss function ℓ such that $0 \leq \ell(h, z) \leq M$, we have,

$$\mathbb{E}_{\mathbf{s}} |\mathcal{L}(h) - \ell_{\text{emp}}(h)|^2 \leq M \bar{\epsilon}(n) + \frac{2M^2}{n}.$$

Proof. Let N_i be the set of index of points of \mathbf{s} that fall into the C_i . Note that $(|N_1|, \dots, |N_K|)$ is an i.i.d. multinomial random variable with parameters n and $(\mu(C_1), \dots, \mu(C_K))$. Then $\mathbb{E}_{\mathbf{s}}|N_k| = n \cdot \mu(C_k)$ for $k = 1, \dots, K$.

$$\begin{aligned}
& \mathbb{E}_{\mathbf{s}} |\mathcal{L}(h) - \ell_{\text{emp}}(h)|^2 \\
&= \mathbb{E}_{\mathbf{s}} \left| \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z) - \frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right|^2 \\
&= \mathbb{E}_{\mathbf{s}} \left| \sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) - \frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right|^2 \\
&= \left(\sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) \right)^2 + \frac{1}{n^2} \mathbb{E}_{\mathbf{s}} \left(\sum_{i=1}^n \ell(h, s_i) \right)^2 \\
&\quad - 2 \left(\sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) \right) \mathbb{E}_{\mathbf{s}} \left(\frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right) \\
&\leq \left(\sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) \right) \left| \sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) - \mathbb{E}_{\mathbf{s}} \frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right| \\
&\quad + \frac{1}{n^2} \mathbb{E}_{\mathbf{s}} \left(\sum_{i=1}^n \ell(h, s_i) \right)^2 - \left(\sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) \right) \mathbb{E}_{\mathbf{s}} \left(\frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right) \\
&\leq M \underbrace{\left| \sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mu(C_k) - \mathbb{E}_{\mathbf{s}} \frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right|}_H + \frac{2M^2}{n}
\end{aligned}$$

We then bound the term H as follows.

$$\begin{aligned}
H &= \left| \sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) \mathbb{E}_{\mathbf{s}} \frac{N_k}{n} - \mathbb{E}_{\mathbf{s}} \frac{1}{n} \sum_{i=1}^n \ell(h, s_i) \right| \\
&\leq \frac{1}{n} \mathbb{E}_{\mathbf{s}} \left| \sum_{k=1}^K \mathbb{E}_{z \in \mathcal{Z}} \ell(h, z | z \in C_k) N_k - \sum_{j \in C_k} \ell(h, s_j) \right| \\
&\leq \frac{1}{n} \mathbb{E}_{\mathbf{s}} \sum_{k=1}^K N_k \max_{s_j \in C_k, z \in C_k} |\ell(h, z) - \ell(h, s_j)| \\
&= \frac{1}{n} \sum_{k=1}^K N_k \mathbb{E}_{\mathbf{s}} \max_{s_j \in C_k, z \in C_k} |\ell(h, z) - \ell(h, s_j)| \\
&\leq \bar{\epsilon}(n).
\end{aligned}$$

Then we have,

$$\mathbb{E}_{\mathbf{s}} |\mathcal{L}(h) - \ell_{\text{emp}}(h)|^2 \leq M \bar{\epsilon}(n) + \frac{2M^2}{n}.$$

□

To analyze the generalization performance of deep learning with dropout, following lemma is central.

Lemma 3 (Bounded difference inequality [25]). *Let $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_L) \in \mathcal{R}$ be L independent random variables (\mathbf{r}_l can be vectors or scalars) with $\mathbf{r}_l \in \{0, 1\}^{m_l}$. Assume that the function $f : \mathcal{R}^L \rightarrow \mathbb{R}$*

satisfies:

$$\sup_{\mathbf{r}^{(l)}, \tilde{\mathbf{r}}^{(l)}} |f(\mathbf{r}^{(l)}) - f(\tilde{\mathbf{r}}^{(l)})| \leq c_l, \forall l = 1, \dots, L,$$

whenever $\mathbf{r}^{(l)}$ and $\tilde{\mathbf{r}}^{(l)}$ differ only in the l -th element. Here, c_l is a nonnegative function of l . Then, for every $\epsilon > 0$,

$$\begin{aligned} \mathbb{P}_{\mathbf{r}} \{f(\mathbf{r}_1, \dots, \mathbf{r}_L) - \mathbb{E}_{\mathbf{r}} f(\mathbf{r}_1, \dots, \mathbf{r}_L) \geq \epsilon\} \\ \leq \exp \left(-2\epsilon^2 / \sum_{l=1}^L c_l^2 \right). \end{aligned}$$

C Proof of Theorem 1

Proof of Theorem 1. Now we proceed to prove Theorem 1. Using Chebyshev's inequality, Lemma 2 leads to the following inequality:

$$\Pr_{\mathbf{s}} \{|\mathcal{L}(h) - \ell_{\text{emp}}(h)| \geq \epsilon|h\} \leq \frac{nM\mathbb{E}_{\mathbf{s}} \max_{s \in \mathbf{s}, z \sim s} |\ell(h, s) - \ell(h, z)| + 2M^2}{n\epsilon^2}.$$

By integrating with respect to h , we can derive the following bound on the generalization error:

$$\Pr_{\mathbf{s}, \mathcal{A}} \{|\mathcal{L}(h) - \ell_{\text{emp}}(h)| \geq \epsilon\} \leq \frac{nM\mathbb{E}_{\mathcal{A}, \mathbf{s}} \max_{s \in \mathbf{s}, z \sim s} |\ell(h, s) - \ell(h, z)| + 2M^2}{n\epsilon^2}.$$

This is equivalent to:

$$|\mathcal{L}(h) - \ell_{\text{emp}}(h)| \leq \sqrt{\frac{nM\bar{\epsilon}(n) + 2M^2}{\delta n}}$$

holds with a probability greater than $1 - \delta$. □

D Proof of Theorem 2

Proof. To simplify the notations, we use $X(h)$ to denote the random variable $\max_{z \sim s} |\ell(h, s) - \ell(h, z)|$. According to the definition of ensemble robustness, we have $\mathbb{E}_{\mathcal{A}} X(h) \leq \epsilon(n)$. Also, the assumption gives $\text{var}[X(h)] \leq \alpha$. According to Chebyshev's inequality, we have,

$$\mathbb{P} \left\{ X(h) \leq \epsilon(n) + \frac{\alpha}{\sqrt{\delta}} \right\} \geq 1 - \delta.$$

Now, we proceed to bound $|\mathcal{L}(h) - \ell_{\text{emp}}(h)|$ for any $h \sim \Delta(\mathcal{H})$ output by $\mathcal{A}_{\mathbf{s}}$.

Following the proof of Lemma 2, we also divide the set \mathcal{Z} into K disjoint set C_1, \dots, C_K and let N_i be the set of index of points in \mathcal{f} that fall into C_i . Then we have,

$$\begin{aligned} & |\mathcal{L}(h) - \ell_{\text{emp}}(h)| \\ & \leq \frac{1}{n} \sum_{i=1}^K \sum_{j \in N_i} \max_{z \in C_i} |\ell(h, s_j) - \ell(h, z)| + \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}} \\ & \leq \epsilon(n) + \frac{\alpha}{\sqrt{\delta}} + \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}} \end{aligned}$$

holds with probability at least $1 - 2\delta$. Let δ be 2δ , we have,

$$|\mathcal{L}(h) - \ell_{\text{emp}}(h)| \leq \epsilon(n) + \frac{\alpha}{\sqrt{2\delta}} + \sqrt{\frac{2K \ln 2 + 2 \ln(1/2\delta)}{n}}$$

holds with probability at least $1 - \delta$. This gives the first inequality in the theorem. The second inequality can be straightforwardly derived from the fact that $\text{var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2 \leq M\mathbb{E}[X] - (\mathbb{E}[X])^2$. □

E Proof of Theorem 3

Proof. Let $R(\mathbf{s}, \mathbf{r}) = \mathcal{L}(\mathcal{A}_{\mathbf{s}, \mathbf{r}}) - \ell_{\text{emp}}(\mathcal{A}_{\mathbf{s}, \mathbf{r}})$ denote the random variable that we are going to bound. For every $\mathbf{r}, \mathbf{t} \in \mathcal{R}^L$, and $L \in \mathbb{N}$, we have

$$\begin{aligned} & |R(\mathbf{s}, \mathbf{r}) - R(\mathbf{s}, \mathbf{t})| \\ &= \left| \mathbb{E}_{z \in \mathcal{Z}} [\ell(\mathcal{A}_{\mathbf{s}, \mathbf{r}}, z) - \ell(\mathcal{A}_{\mathbf{s}, \mathbf{t}}, z)] - \frac{1}{n} \sum_{i=1}^n (\ell(\mathcal{A}_{\mathbf{s}, \mathbf{r}}, z_i) - \ell(\mathcal{A}_{\mathbf{s}, \mathbf{t}}, z_i)) \right| \\ &\leq \mathbb{E}_{z \in \mathcal{Z}} |\ell(\mathcal{A}_{\mathbf{s}, \mathbf{r}}, z) - \ell(\mathcal{A}_{\mathbf{s}, \mathbf{t}}, z)| + \frac{1}{n} \sum_{i=1}^n |\ell(\mathcal{A}_{\mathbf{s}, \mathbf{r}}, z_i) - \ell(\mathcal{A}_{\mathbf{s}, \mathbf{t}}, z_i)|. \end{aligned}$$

According to the definition of β :

$$\sup_{\mathbf{r}, \mathbf{t}} |R(\mathbf{s}, \mathbf{r}) - R(\mathbf{s}, \mathbf{t})| \leq 2\beta,$$

and applying Lemma 3 we obtain (note that \mathbf{s} is independent of \mathbf{r})

$$\mathbb{P}_{\mathbf{r}} \{R(\mathbf{s}, \mathbf{r}) - \mathbb{E}_{\mathbf{r}} R(\mathbf{s}, \mathbf{r}) \geq \epsilon | \mathbf{s}\} \leq \exp\left(\frac{-\epsilon^2}{2L\beta^2}\right).$$

We also have

$$\mathbb{E}_{\mathbf{s}} \mathbb{P}_{\mathbf{r}} \{R(\mathbf{s}, \mathbf{r}) - \mathbb{E}_{\mathbf{r}} R(\mathbf{s}, \mathbf{r}) \geq \epsilon\} = \mathbb{E}_{\mathbf{s}} \mathbb{P}_{\mathbf{r}} \{R(\mathbf{s}, \mathbf{r}) - \mathbb{E}_{\mathbf{r}} R(\mathbf{s}, \mathbf{r}) \geq \epsilon | \mathbf{s}\} \leq \exp\left(\frac{-\epsilon^2}{2L\beta^2}\right).$$

Setting the r.h.s. equal to δ and writing ϵ as a function of δ , we have that with probability at least $1 - \delta$ w.r.t. the random sampling of \mathbf{s} and \mathbf{r} :

$$R(\mathbf{s}, \mathbf{r}) - \mathbb{E}_{\mathbf{r}} R(\mathbf{s}, \mathbf{r}) \leq \beta \sqrt{2L \log(1/\delta)}.$$

Then according to Lemma 1:

$$\mathbb{E}_{\mathbf{r}} R(\mathbf{s}, \mathbf{r}) \leq \bar{\epsilon}(n) + \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}}$$

holds with probability greater than $1 - \delta$. Observe that the above two inequalities hold simultaneously with probability at least $1 - 2\delta$. Combining those inequalities and setting $\delta = \delta/2$ gives

$$R(\mathbf{s}, \mathbf{r}) \leq \beta \sqrt{2L \log(1/\delta)} + \bar{\epsilon}(n) + \sqrt{\frac{2K \ln 2 + 2 \ln(2/\delta)}{n}}.$$

□